An Interview with

MARVIN L. MINSKY

OH 179

Conducted by Arthur L. Norberg

on

1 November 1989

Cambridge, MA

Charles Babbage Institute
Center for the History of Information Processing
University of Minnesota, Minneapolis

Marvin Minsky Interview
1 November 1989

Abstract

Minsky describes artificial intelligence (AI) research at the Massachusetts Institute of Technology (MIT). Topics include: the work of John McCarthy; changes in the MIT research laboratories with the advent of Project MAC; research in the areas of expert systems, graphics, word processing, and time-sharing; variations in the Advanced Research Projects Agency (ARPA) attitude toward AI with changes in directorship; and the role of ARPA in AI

research.

MARVIN L. MINSKY INTERVIEW

DATE: 1 November 1989 INTERVIEWER: Arthur L. Norberg

LOCATION: Cambridge, MA

MINSKY: It would be fun to trace wordprocessing, because the first wordprocessors were the editors, program editors, that came with time-sharing, so that we had a little one... There were two of them. There was one called Expensive Typewriter and I think John McCarthy or someone in his group made that just for editing programs, since you didn't have to wait a day for your program to come back. You wanted to type it right into the machine. Of course, the time-shared machines were the first ones that let a number of people operate a computer from the terminal without huge cost. The earlier machines usually had a typewriter, but only one person could work it. And then, I think Ed Fredkin wrote another one called Colossal Typewriter, which wasn't very big either. We started using that to edit text as well. The big bottleneck was that all the printers were things like teletypes; only a few of them had upper and lower case.

NORBERG: When was that problem solved in moving from fully upper case to some combination?

MINSKY: That's interesting. There was a machine called the Flexowriter, which had upper and lower case and could be connected to a computer, so most people used those. After the beginning of Project MAC, in 1963, a bunch of us were really fed up because we had basically teletype terminals. So, Fano and me and I think, Licklider, I'm not sure who else went, we all went out on this trip to Skokie, Illinois, to visit Western Union... No, the Teletype Corporation to try to convince them that there was a huge market for upper and lower case machines. They didn't really believe us much, but they had one called the model 33, which was not for sale yet. Maybe this influenced them to get it out, because within a couple of years we got the model 33 teletype which had both cases. But for better or for worse, they blew it. Teletype just doesn't appear in the automatic typewriter market for some reason. Maybe because they made machines that were so rugged and reliable. I remember on the visit we saw some Teletype working under water; they tested them under all sorts of conditions. So, they probably couldn't afford to make a consumer type product.

NORBERG: Do you remember when that was? Was that early in Project MAC?

MINSKY: Yes. I'm pretty sure it was the first year or two, '63 or '64, maybe Fano would remember.

NORBERG: I was just curious, because Alan Kay mentioned Flexowriter the other day and mentioned it for a later date, which is why I was....

MINSKY: No, Flexowriter I think had been around. I did the manuscript of "Steps Toward Artificial Intelligence" with a Flexowriter, I think.

NORBERG: Oh, I see, '61 then.

MINSKY: It was a big pain, but it was better than retyping. And it wasn't on a computer. I just had one home and I would run the tape through and type it out and punch more holes.

NORBERG: I've read most of the works that have appeared about you and I've read a number of things that you've written. So I'm not going to spend any time asking you personal questions or trying to do what Bernstein did for the *New Yorker* magazine or any of that. I would like to focus on two issues in this interview. One of them is - as you could tell me better, I'm sure - it's very difficult to decide where to enter the question of what is AI, and what sorts of sub-areas of AI I should focus on in looking at the history of development because we obviously can't do the history of AI; it's too big a field.

MINSKY: Well, I think the problem is that the things that we call the AI labs socially were really... nobody knew what AI was. I have three definitions.

NORBERG: That's exactly where I'd like to start. I'd like to ask you if you can review for me the determination of AI objectives in the late 1950s.

MINSKY: OK. The simplest one is: how could we build a machine that thinks and does all the sorts of things that you would say is thinking if a person did them? So that's one question. Another question is: is there a theory for this? Could we understand principles of intelligence or something like that? Those are not quite the same thing, because somebody might think maybe there are some general principles of problem solving that apply to any machine that can solve hard problems, whether it's a person or not. But the third one and the one that's really the most effective definition is very weak, namely the AI labs are the places where young people go if they want to make machines do things that they don't do yet and that are hard to do with the currently existing methods. So in that sense, AI is just the more forward looking part of computer science. That's the definition that makes sense in terms of the history, because if you go to the AI labs you'll see people doing... At one time they were doing advanced graphics of one sort or another and then that becomes a little industry of its own, or they're working on speech recognition and that becomes a separate thing. From the very start AI labs were obsessed with making machines that could see. So there was the early vision work. A lot of the early vision work was somehow associated with AI. Although my friend Russell Kirsch at the Bureau of Standards, had been doing it as early as 1954 or '55.

NORBERG: Which? Vision?

MINSKY: A little bit of computer vision with the SEAC machine. Nobody used the word AI yet. He was the kind of person who could say well, everybody's having their machines crunching numbers, can we have one that can recognize patterns. Of course, the electrical engineers would dispute that and say, we've been making sensory systems all the way back to World War II. I think if you could get a portrait of an AI lab at any particular time you'd see that only a few of the people are trying to make a thinking machine and other people are just trying far out, new kinds of software.

NORBERG: Is this a definition that's imposed on this after the fact, or was it clear in those days of the late 1950s that this is what people were doing?

MINSKY: Well, there were... Certainly when we started the MIT AI Lab, the goal was let's see if we can make a machine that will have common sense and solve the kinds of problems that people do. Now, Newell and Simon in their Carnegie-Mellon and Rand places had similar ideas. The main goal of their work was to imitate thought; for some reason they decided to call it complex information processing. They didn't like the term AI really, but I think that was sort of justification for the labs and the way they got funded and the way they explained it to people. But this other part of just being in the forefront of new computer techniques and hackery and getting the machines to do all kinds of new things was just the informal thing that actually happened.

NORBERG: How did support come for these groups? Let's just stick with your group for the moment.

MINSKY: Okay. Our group was... It was very simple from my point of view, McCarthy and I met - he might have a slightly different version of it, we were walking down the hall and we met Jerry Wiesner or Zimmerman or someone and he said how's it going and we said well, we're working on these artificial intelligence ideas but we need a little more room and support for some graduate students. So then a room appeared a few days later, I believe because it was under the auspices of RLE, the Research Laboratory Electronics had a block grant, so did Wiesner and/or Zimmerman, I don't remember quite who, could simply decide let's give them a room and some assistance.

NORBERG: Do you know where the block grant was from?

MINSKY: It was some joint service thing. I don't know exactly what the original service was. Was there a joint services?

NORBERG: There was. The Tri-Service Program.

MINSKY: I'll bet it was that, but I can't be sure. Also, then the computation center at MIT started, I think '61 but I'm not sure. Philip Morse was head of that, and I think that had money from IBM.

NORBERG: It certainly had equipment from IBM.

MINSKY: It certainly had equipment and I bet there was some money. As I recall, that paid for some... we had three or four graduate students and I think that paid for that. So, my recollection is that all this was done without ever writing a proposal.

NORBERG: That's what I was just about to say. You people didn't have to be concerned at all with money.

MINSKY: Zero. Right. We had people like Wiesner and Morse, who just did it. Then when it was starting to get large and then the miracle of Project MAC happened because just at the time that we were starting to need a lot of money for our own machines, '62 - '63, then the Project MAC miracle happened and Project MAC got 3 million dollars a year and we got 1 million dollars a year. I don't recall ever writing a proposal for that.

NORBERG: But in the meantime, quite a number of memos came out from the AI project which was in RLE and the Computation Center.

MINSKY: Yes. We started this idea that whenever anybody got a good idea, they should write it out. There wasn't any idea of wide publication but the memos got numbered and there got to be several hundred of them before long. Memos three through eight are missing; nobody has a copy, but all the rest we....

NORBERG: Yes, managed to recover. Well, how did this interaction with Fano then go on?

MINSKY: It was just incredible. We just managed everything sort of informally. I think the unique thing about our lab was that it was co-directed. McCarthy and I simply... Whoever was in would make whatever decision had to be made. Then McCarthy went to Stanford and Papert came to MIT and we continued that.

NORBERG: How did this interaction between you and McCarthy go on? This seems rather haphazard, if you'll pardon the word.

MINSKY: Oh, yes. Well, we had known each other quite well in graduate school in '53 and '54, or so, '52.

NORBERG: What sort of relationship was that? Did you work on problems together, did you take classes together, dissertations closely allied?

MINSKY: No, dissertations were pretty different, but it was a social thing. First of all the mathematics department was rather small; everybody knew each other and everyone had many common interests. I was a bit interested in topology. Differential topology was his main concern. There was another mathematician named Nash that I worked more closely with. I never really worked with McCarthy on anything mathematically. There was Lloyd Shapley. Shapley and Nash were the key people in the Rand Corporation game theory business and Martin Schubeck was an economist, also a mathematician. The whole cluster of these people all worked on different things and talked about them to each other.

NORBERG: But, now, how does this then develop into a relationship of the kind that you're describing to me?

MINSKY: We used to talk about how you could make thinking machines. I had been working on it longer than anyone else because I was enamored of Warren McCullough and people like that in this little college. And then I went to Harvard at the end of graduate school. I went to Tufts for six months - an ill-fated, ambitious project - and then I got the junior fellowship at Harvard. McCarthy went off to Dartmouth, which was sort of a nice miracle there.

You know what happened? You could say, how come Dartmouth was prominent in computer science? What happened was that, as I understand it, the four professors all retired about the same year.

NORBERG: Four people in mathematics?

MINSKY: At Dartmouth. So, Dartmouth called up and said what do we do? Where's the best mathematics department? They called up Princeton and asked someone, perhaps Lefschetz who was the Dean, I don't know, we need mathematicians, and they got four of them. Kemeny came as, I don't know if he was chairman yet, but he had finished his degree a year or two earlier than the others, and then Carl Delieu (sp?), I think, I'm sure Snell and McCarthy -I forget who the fourth one was - it wasn't Delieu. But they got four young mathematicians all at once, and McCarthy was already interested in the possibilities of computers. Then we didn't see each other very much for several years, but we got together again in '56 - well, just two years - and McCarthy organized that AI conference. We kept in closer touch then, and my fellowship got ditched in '57 and I went to Lincoln in '58, MIT in '59. McCarthy came down there the year before and we decided to start a group because we were both working on common sense reasoning.

NORBERG: All right, let me leave that for a moment and talk about Lincoln because there are two things about that that interest me that have not shown up in any of your other work, so it's hard for me to know exactly what their interests were in common sense reasoning.

MINSKY: It was Oliver Selfridge. There was one little group there called Group 34 and Oliver Selfridge was the leader of it. He was building a machine for the Air Force, I think, but I'm not sure, to read Morse code. Nobody had ever made a machine that could read hand-keyed Morse code. It was quite important for security reasons as well as communication reasons because these people were still sending uncrypted messages from all over the world and it would be nice to have a machine that would just listen. It's quite hard to do, and in order to do it he had developed quite an advanced theory [unintelligible] a pandemonium of having a lot of different little detectors and other ones listening to them and so forth.

NORBERG: And did this machine get built and work effectively?

MINSKY: I think it worked semi-effectively; I don't know if it was ever actually used. It was fairly big and expensive because a flip-flop is this big. We got annoyed at the size of the flip-flops; a couple of guys there made smaller ones in little modules. They turned out to be Ken Olson and Ben Gurley, who started a company to sell these modules. First they had everybody at Lincoln wanting them. I had about 30 of them myself for this bridge (?) detector I was building.

NORBERG: I see. What was your role in Selfridge's project?

MINSKY: I was just a staff member working on my own things.

NORBERG: A staff member working on your own things. How can you be a staff member on that project and working on your own things?

MINSKY: Oliver had worked for McCullough, so he came from MIT. McCullough's lab usually had five or six people usually developing their own theories and Oliver started this Lincoln group. He didn't fill the slots with people who would do jobs; he got the smartest people he could find and they all managed to work together and develop these things. Brad Howland was another one who was interested in mechanical electronics and optics and things. So when I got there, Oliver just let me do whatever I wanted. We had actually been working together anyway. We developed some theory of perceptrons and linear learning networks and gradient optimizing devices, a whole lot of things that came out of cybernetic work.

NORBERG: Was that before you went there or while you were there?

MINSKY: We did a lot while we were there, but we had done some in the previous year or so.

NORBERG: So you stayed there for two years and then went off to MIT?

MINSKY: Yes, I think it was just one year and then I got an invitation from Ted Martin to be in the math department.

NORBERG: And to pursue your own research obviously?

MINSKY: Yes.

NORBERG: What did you plan on doing once you arrived at MIT? Did you re-establish this relationship with McCarthy right away?

MINSKY: Yes.

NORBERG: And the two of you set out to do what?

MINSKY: To get good students and give them problems. McCarthy already had this LISP system under development, so there were quite a few students who were helping build that.

NORBERG: I was looking at some of those AI memos today, as a matter of fact.

MINSKY: A lot of the early ones were just details about the LISP system.

NORBERG: Yes, I was trying to decide when your first one appeared, and it seems to me it's quite late actually in comparison to some of the others. Memo 17, which was the first one I took notice of, it's the one on problems with common sense which was done by McCarthy, and that's November 1958. From that point on they appear rather profusely.

MINSKY: Yes, so he must have come to MIT in '58, a year before I did, I'm not sure.

NORBERG: How aware were you at this time of the work that was going on elsewhere? Such as at Carnegie Institute of Technology then, at SRI, maybe other AI work that was going on that I don't know about.

MINSKY: I was very, very current on the Carnegie-Mellon stuff - it was the best stuff being done. I don't remember when I established connections with SRI.

NORBERG: Is it possible that was after McCarthy went out to Stanford?

MINSKY: No idea. You could ask Charlie Rosen. Well, let's see, it must have been before '64 because... maybe it was around '64. What happened there was that they were working on perceptrons and I must have known about that because I knew pretty much about all the neural net work going on and I was beginning to -- Seymour and I working discovered the limitations of perceptrons so we must have gone... No use reconstructing history that way. (Laugh) I know I went out there and looked at their machine and talked to Rosen about the future of such things. The key thing would be this: at some point he hired Bertram Raphael, who was one of my graduate students, and the idea was that Raphael would go there and teach them about symbolic processing. So, if we can find out when Raphael got this job, then we know that I had been talking to them a year or two before that.

NORBERG: I want to go back to the idea of objectives again and how one decides what AI is in this period. The three definitions you gave, if we choose any one of them, it seems to me, are fairly broad definitions, and at any time in a new field the broad definition provides some indication of what sort of problems one might work on, but with a whole range of problems possible to work on, how does one decide which ones to address first? How did you people decide?

MINSKY: Well, first, we didn't do the same thing. John decided that a way to understand common sense would be to try to make a mathematical formulation of it, and so he went in the direction to use mathematical logic notations. I think calculus was the most well-established tool, and most of his students went along in that direction. I didn't think that that had much future, still don't. So my students worked more on trying to make machines solve new kinds of problems. So I had one student (Henry Ernst) working on how to make a mechanical robot find objects on a table; another student, Jim Slagle worked on a machine that could do it, solve integral calculus problems. Manuel Bloom is another student who worked on the theory of ultimate computability, which didn't have much to do with AI, but I had also been interested in. Both McCarthy and I shared another interest which was the foundations of computation and the basic mathematical principles of computers. Again, we worked in somewhat different directions. I worked on the theory of Turing machines and Manuel worked on... his thesis was the first one to start developing this idea of computational complexity: when could you make one program twice as fast as another and at what cost. But then, the big difference was that I started working with a student named Tom Evans to make a machine reason by analogy, and that was very different from what the logic people were doing. Dan Bobrow working on making a machine understand algebra problems in English. So I was going more in the heuristic direction of taking problems that people could do and just picking out interesting ones and trying to get machines to do them.

NORBERG: Where were the various heuristic methods coming from? Were you people devising these yourselves, basically?

MINSKY: Yes. These were coming from us and from Carnegie-Mellon.

NORBERG: Were those methods used elsewhere?

MINSKY: No. Pretty much people... You see the basic difference between AI programs and other ones, and it's not really very profound, was that in typical computer programs you would have to say in advance what should be done and in the AI programs you'd say, well, if you're stuck try different things, generate a list of things to try. So, that was the basic idea of search which I think is sort of the characteristic difference between AI and other things. But then the searches got too large so you'd invent heuristics. One kind of heuristic would invent new things to search; another one would discard them. Generally, discarding was more important.

NORBERG: Yes. When I think of something like means/ends analysis, for example, I think of philosophers using that for a good long time before mathematicians picked it up. I'm trying to decide whether you people were out there sort of looking for these or inventing these techniques.

MINSKY: Well, that's a common sense idea. The Newell/Simon thing was basically given a goal and a situation look at the differences and then you'd have a smart machine if you could tell it how to get rid of each difference. So that's a sort of look at the difference between where you are and where you want to be and doing something. You could regard that as means/ends. Nobody had ever made machines do that.

NORBERG: At MIT then there were already the makings of a branching: one into formal logic techniques for this and another to making machines learn somehow.

MINSKY: Yes, for a computer heuristic search. In fact learning came much later. The first learning machines didn't seem to work very well, and so both McCarthy and I decided that we'd solve that problem some day, but you probably couldn't get a machine to learn something unless it had some way of representing what it was that it would learn, so that was perhaps the big difference. The people working on perceptrons and neural nets never agreed with that. They said, well, the machine can learn how to represent things while it's learning. I should also say that the book of Polya was pretty influential. He was a mathematician who wrote a book full of ideas about how people solve problems and described procedures that were probably wrong mostly, but the heuristics came from there.

NORBERG: I see. One incidental question that occurred to me as you were describing what some of your students were doing; how limited were you in those early years by the machines that were available?

TAPE 1/SIDE 2

NORBERG: Was it just the memories were too small or were there other problems of accessibility and reliability as well?

MINSKY: Well see, the time-sharing project came out of the AI things because of the accessibility problems. So, I suppose you could say we felt very badly limited by accessibility.

NORBERG: Came out of the AI?

MINSKY: Yes.

NORBERG: Could you be more explicit about that for me?

MINSKY: McCarthy developed these things, first at Bolt, Beranek and Newman and then at MIT so that we could debug our programs better.

NORBERG: So, just because McCarthy did this makes it AI or because it was associated somehow with an AI problem?

MINSKY: I never used the word AI. That's somebody else trying to structure this rather integrated project of trying to get machines to do more. See, if you can take the first definition of AI, then time-sharing is AI then, because nobody had done it very well. Developing time-sharing had the same role in our minds as developing LISP, mainly how were we going to get our programs to be developed and run.

NORBERG: I guess what that makes me think is that you would classify any new area in computing as on the frontier and therefore somehow associated with what we now, people like me, associate with AI.

MINSKY: Well, that was my first definition. The AI labs were filled with the kind of people who wanted to do new things, and now there were such people in other labs. But they were handicapped because there was only a couple of them in each lab, but if you looked in the MIT lab we had a dozen people, each of whom was trying to start a new domain of computation. People whose names aren't very well known yet - Greenblatt and Holloway and Tom Knight, William Gospers, each of them were pushing in some direction on this frontier. So it was an atmosphere. I got it from McCullough's lab, because there they were interested in the nervous system, but they were inventing new algorithms before computers. Those were the leaders of cybernetics before AI.

NORBERG: What were other people doing that we would call computing in that period?

MINSKY: Mostly ordinary science as Kahn puts it.

NORBERG: Such as?

MINSKY: Well Fortran was the language that had been successful and so those places were making new versions of Fortran called Algol or... lots of names, Pascal. All the same as far as I'm concerned.

NORBERG: So any of these new high-level languages can be considered as just ordinary science?

MINSKY: They were all Algol to me, Algol 1.7, Fortran 2.3 and it's still that way in other places, and rightly so. (Laugh)

NORBERG: How did you first become involved with Project MAC, do you remember?

MINSKY: I don't quite remember how it worked, but Fano says it was on a train ride. He and I and maybe Licklider talking about what should we do? Especially since Licklider was in the position to do it.

NORBERG: Do you have any recollection of that at all other than Fano telling you this after?

MINSKY: I have just a little... No, hardly any clear idea.

NORBERG: Okay.

MINSKY: I never planned more than a few minutes ahead consciously in my whole life.

NORBERG: Let me try to pull it out in a slightly different way. Do you remember in the early '60s various meetings that you went to?

MINSKY: Yes. I remember one which was in some great place in Virginia where we went skeet shooting. (laugh)

NORBERG: You remember the skeet shooting, but do you remember what went on at the meeting?

MINSKY: Not in the least. Hardly anything...

NORBERG: Not at all? Because there was such a meeting held in .... I want to say Hastings, Virgina but it's not Hastings, but it's around the Wye Plantation, in which the Air Force apparently sponsored the meeting. There was some discussion about the directions in which the Air Force program and various computing programs of the military services and the Defense Department ought to go. It was on a return trip where there was some discussion on a train about what should happen.

MINSKY: Yes, right. I presume that we decided those people were going nowhere, what should we really do, and for the first time we were in some place to spend actual hours together instead of meeting in the hall.

NORBERG: I don't remember anyone ever before saying that you were on that train ride.

MINSKY: Fano says I was.

NORBERG: Fano says no, but no one ever... Okay, I don't really want to press that, because I don't think it's useful to press it.

MINSKY: I think I was. I do have an image of it. But one problem is that I had spent an awful lot of time with Licklider, because I used to consult at BBN. Both McCarthy and I and lots of our students used to go over there. You know why? They had a PDP-1 and we didn't.

NORBERG: Can you tell me something about that interaction with Licklider, then? When did you first come into contact with Licklider?

MINSKY: I don't know. Oh! Of course I know - as an undergraduate in 1948 it must have been. I took a course called Physiological Psychology. I think that's in the Bernstein article.

NORBERG: It is.

MINSKY: I had been thinking about learning machines and that sort of thing, but I was sort of all by myself going to Widner Library and reading books by Ruchevsky (?) and by McCullough, and all those things.

NORBERG: How did you find these articles and the book by Rushevsky? Were these discussed in any program you were involved in?

MINSKY: No. I don't remember how I started reading them. Oh! Of course, I had read the books, there were these books coming out in the late '40s called cybernetics - the Macy conference volumes, and I read those word for word and got really excited. McCullough was one of the people [unintelligible] so that must have been it. Rushevsky was in one of those conferences and I happened to find Rushevsky's book, which was quite big and each chapter was another mathematical model of how a cell might divide or something like that. It was very nice to see that you could do things like that because if I did them myself I felt a little bit like a crackpot. (Laugh) And Licklider's course was a revelation because there were graduate students and they didn't understand anything and I actually wrote some memos in that course for him. He kept saying that's great. George Miller was also down there in the laboratory. I remember spending a lot of time with George and explaining to him the McCullough/Pitts paper which was full of errors. All those people thought it was their fault; in fact Pitts must have been bluffing. He didn't quite have a theory ready.

NORBERG: I see.

MINSKY: All the psychologists felt it was something wrong with them.

NORBERG: Did you keep up this contact with Licklider or did it reoccur later on?

MINSKY: I think it reoccurred.

NORBERG: Do you remember when?

MINSKY: I kept contact with Miller all along. He was still at Harvard when I was a junior fellow.

NORBERG: What sort of projects did you consult on out at BBN?

MINSKY: I had one student write a program for teaching children algebra. And I think that was the precursor to Slagle's thesis.

NORBERG: Do you remember who that student was and what the program was?

MINSKY: His name is Steve Savatsky.

NORBERG: Doesn't mean anything to me.

MINSKY: Disappeared. But the program would just do simple things like you type in an equation and ... Oh, it taught you about parabolas and would ask you what's the lowest point. You would type that and it would say that's impossible because a = 3. (laugh) Some dumb little interactive thing. It would probably be fairly good now.

NORBERG: What was BBN's interest in that?

MINSKY: I don't know how they pulled this off, but you see what they had at BBN was making money on acoustics and there was a miracle in there. That place was like a university psych department. I'd come over and it would remind me of what it was like in the basement of Memorial Hall. They had different people working on different theories of perception. There was a sort of close bridge to the acoustics because somebody would be working on, how a subject perceived how loud a sound is, and, let me see, it was pioneering in cognitive things. Dan Bobrow and other students would work on other things. I think it was getting funded by AFOSR or something, so BBN didn't care, since the project was paying for itself anyhow.

NORBERG: What sort of opportunities did this provide for your students?

MINSKY: It meant that students could go there and earn some money while they were pursuing their studies and still do the same thing.

NORBERG: Which was fairly traditional at MIT, it seems to me. I remember in the '20s and '30s that that was going on with professors having some sort consulting relationship with the various companies - Raytheon is one that comes to mind.

MINSKY: That's right. Now, I think at this time Licklider was full-time there. He had been at MIT for a while, I think, then he was at BBN and he didn't come back to MIT until later in Project MAC.

NORBERG: After he left IBM, after one or two years.

MINSKY: I didn't know he went to IBM.

NORBERG: He went to IBM from DARPA and spent almost two years there before coming back. So you had this relationship. We got onto that because we were talking about the train ride and you becoming involved in Project MAC. With this new amount of money...

MINSKY: You see, you're putting things together for me again because part of the chemistry that made that work was probably that Licklider and I had a great deal of confidence and respect from my college days. And he was just an assistant professor then himself, so there wasn't that much difference. I don't know, was he 10 years older?

NORBERG: You mean back in '48?

MINSKY: Yes. So that we knew each other a lot.

NORBERG: With all this new amount of money, this million dollars now that's available ostensibly for AI in 1963, how did that change your outlook or vision for the research laboratory?

MINSKY: Oh good. That hits what we were talking about in the proposal thing. What was the limitation of AI research? In '58 or '59, the limitation was that the 704, the IBM machine, was more or less fast enough for what we wanted to do, but it didn't have enough memory. By 1962, the programs were starting to occupy two coreloads, which means you would have to run something and then re-load and run something else and so 32,000 words of memory were beginning to be a pain. Our staff was getting sort of slowly larger; it didn't expand dramatically. Mostly the people we had in the project in the late '60s were undergraduates who had grown up in it and stayed with it. That's another maybe unusual feature. So the million dollars came in very smoothly because we spent most of it the first year on buying, I forget what, but we got our own PDP-1, which was a few hundred k with all the stuff, and then before long we got a PDP-6, which was DEC's next big machine. My impression is that you might look at that money, which lasted quite a few years, a million dollars a year, and the first few years it was spent mostly on equipment as we needed it, and then more on people. So although it sounds like a huge influx of money, it was not used for ultra-rapid expansion, but it was very luxurious. If there was somebody we were interested in, we could invite them to come for a year, absolutely. There was hardly ever a chance of deciding should we do this or that; we could do both. This was all without any planning, but the hardware got less expensive as the people got more expensive.

NORBERG: How did the new hardware change the nature of the problems that were being worked on?

MINSKY: It just meant that you could write bigger programs. I think Joel Moses and Bill Martin started writing MACSYMA in '65 or '66 and those required the biggest machine available because MACSYMA was a huge LISP program that would hardly fit. What we did then was we went and spent another million dollars to buy more memory for the PDP-6. So we negotiated a contract with a company called Fabri-Tek and got the largest memory in existence except secret ones.

NORBERG: Fabri-TeK, wasn't that a Minnesota firm at the time?

MINSKY: I think so, yes.

NORBERG: Subsequently bought up by a company in Indiana, as I remember.

MINSKY: We got them to build this huge core memory with about 10 million bits in it or something. They had never made anything that large.

NORBERG: What was it attached to - a 7090 or a PDP-6?

MINSKY: PDP-6. And then, you see, we had these students who had grown up... Stuart Nelson is well-known, at least for the [unintelligible] automatic paging... [unintelligible] So it might have been the most advanced computer around - all time-shared.

NORBERG: How many people using it at the same time?

MINSKY: Usually about 20.

NORBERG: How was the field evolving at that point? Were the objectives changing in any way, considering the new interest on the part of the Defense Department in this research?

MINSKY: No, I don't think so. There weren't any sudden transitions until about '72 or '73. What evolved in this period was that the emphasis moved from... See, at the end of the '50s the problem was we knew how to solve some problems like proving theorems by making a big search and we needed heuristics to make the search smaller. And then, at least in my side of the thing, we began to sense that heuristics would not go very far really when you had a big tree. So, the emphasis shifted to saying, in order to solve a problem the machine has to know, has to have built-in knowledge of what's relative to that field, so we gradually turned to finding new ways to represent knowledge and building up these knowledge bases and that led to these kind of expert systems.

NORBERG: And you said that transition is in the early '70s?

MINSKY: Well the transition of saying we've got to get the machine to have the knowledge for doing these things, that's in the middle sixties.

NORBERG: It is? That sounds early to me.

MINSKY: Yes. Well, the thing like MACSYMA had virtually no heuristics at all. We started to get knowledge about algebra and rules about when to apply them, so it did very little searching.

NORBERG: What was the meaning in your comment when you said, until the transition in the early '70s? What transition were you referring to?

MINSKY: Oh, that was when the Mansfield Amendment and things started, and we had to start writing complicated proposals to ARPA justifying what we were doing and had to find some relevance, things like that. It got to be rather painful. That wasn't in the research itself. Eventually Papert and I quit, because we couldn't stand that, having to be able to manage it.

NORBERG: What sort of justifications were given before that?

MINSKY: Didn't have to. Before that, people like Larry Roberts and Ivan Sutherland and Bob Taylor were the interface, and they would explain to their superiors why this was relevant and they never bothered us at all. You see, in the late '60s when... when did Taylor leave?

NORBERG: '69.

MINSKY: So then he went to BBN - I shouldn't say BBN right? (laugh) Because Xerox was BBN to us. Dan Bobrow went there. Licklider had left so Xerox was the new BBN and Bob Taylor went to BBN west as far as I was concerned. But at about that time, then the pressure was starting to come from ARPA that IPTO couldn't insulate us from. We would actually meet with Fubini and people about how to reformulate a proposal so it would be acceptable. And that's when we both quit.

NORBERG: Do you remember any examples of this? As to how you reformulated it to make it acceptable to the Defense Department?

MINSKY: As I recall, the compromises were almost invisible, so I can't remember any specific things. Sometimes we'd say well, this vision system would be usable to spot cats in the bushes, but I don't think I ever wrote any of those. But it just... then we would have to make milestones... Actually it wasn't so much that it was military as that it was bureaucratic. We'd have to say, we hope to get this vision system doing this by October. So, that's very different from... But I sort of attribute it to the idea that they were having more Congressional oversight and having to cover their tracks and in the early '60s that just never reached us.

NORBERG: Who did have to be satisfied about the research?

MINSKY: Beats me. (Laugh) You know, I've met the directors of ARPA since then and they always said they had perfect confidence in me.

NORBERG: Well, I guess I wasn't referring to the ARPA people, but how does one get validation for one's research in this new area of AI in the 1960s?

MINSKY: You mean, as opposed to whether the proposal is acceptable?

NORBERG: Yes.

MINSKY: That's a funny question; it almost never came up. It's not like physics where... In physics there's a very strong culture of what problems have to be solved. You have to explain why the meson weighs three hundred or whatever it does and everybody knows that if you have a theory that comes out - I gather there still isn't any decent theory for that. (Laugh)

NORBERG: As far as I know there isn't, but I don't keep up with it anymore.

MINSKY: But in AI, if something was hard you could do something else, and...

NORBERG: Just simply abandon the problem?

MINSKY: Well, you just leave it cooking because there's no point in... I think the thing is that physics had only one problem; we've got to find the unified field theory or the five laws of nature or something. Now, it's true that that's not true. Namely, you can also find a stronger alloy for making piano wires or infinite practical things. But, right at the heart of - if you say you're doing theoretical physics - you don't have a lot of choices on what to do. But if you're doing intellectual things, you could say, well it looks like Greenblatt's coming along very well on the chess, program let's go on with that, but the vision system just doesn't seem to be getting anywhere, so we'll do more of MACSYMA. The point is as far as validation for the world is concerned, it was these Ph.D. theses mainly and when Tom Evans' analogy program worked everybody was quite impressed and when Bobrow got the thing to do teaching high school algebra a little bit, people would say well, that's pretty impressive. They didn't have anything to compare it to. Now you would look back at it and say, look, it parsed these pieces of English, but if you had this kind of clause it crashed. But since it was the first one that did any kind of clause for... See, there were other parsing programs, at Harvard especially, which were parsing English, but none of the other people were taking this English and understanding what it said and solving the problem.

NORBERG: Which programs were doing the parsing at Harvard?

MINSKY: Oh, I think Hobinger had some and Kuno had some and...

NORBERG: So, it would be Synthex and programs like that?

MINSKY: Yes, three or four of these syntactical things. But I think there was one, Stone's thing called the General Inquirer, which didn't have much input. It actually solved things by answering questions. As far as I know we were the only ones making things that would try to understand the English a little bit. So in a sense, it really is a pathological piece of science where nobody's done anything much yet. The computers have only been available since... Scientists only got computers in the late '50s. There were the SEACs and the ENIACs and things. It's only in

1957 when Fortran comes out that there are actually hundreds of people who get to use them. So, everything we did was new and you didn't have this question of even why does it really work.

NORBERG: That suggests to me, though, that there's no peer review process going on within this avant garde area of computer science.

MINSKY: There's no formal one; there was an internal one.

NORBERG: How did the internal mechanism work?

MINSKY: The internal mechanism would be somebody did something that she did by having the knowledge right near the surface, then everybody would consider them a loser...
(Laugh)

NORBERG: What does it mean to have the knowledge right near the surface?

MINSKY: I mean you could have a question answerer where the answers are in there and it's just trivial. (Laugh)

NORBERG: I see.

MINSKY: I don't think there was much of a problem of... The peer review was very... Oh! here's another. This is sociology. The thing about AI was that nobody believed it could be done. Our worst enemies were other computer scientists. A lot of people thought that it would be humanists and artists who were skeptical, but actually it was always people down the hall.

NORBERG: What was their fear or their concern?

MINSKY: Their idea was it's impossible. A computer only does what it's programmed to. You can't solve problems unless you cheated by dumping the solutions in and there's just... I mean computers don't have souls. It's a religious idea and it's still there a bit. But the point is that it meant that the AI people were a sort of beleaguered minority. I think people like Newell and Simon and us and people at Stanford and SRI as forming a kind of circle. So, the peer view doesn't exist in the other sense, but there's a sense that what we do has to be very good or these vultures will get us. One great feature was that if you discovered something you'd call them up the same day and say, hey I found this way to... Why don't you make GPS work on this kind of difference instead of that.

NORBERG: This is calling up Newell or Simon or one of those people, not the Neanderthals that are not paying attention to you.

MINSKY: That's right. So, we'd always be on the phone and I think there were no secrets. That meant if you had a bad idea, you'd probably get peer reviewed in a minute instead of in a journal a year later. Now that's changed a bit. Not so much that the skepticism has gone, although most of it's gone, but because of this damn commercial problem that... Like, I have a student who's made an optical score reader, reads piano scores rather well, and somebody else somewhere else might have one, so they might say, I don't want the details of this to get out. When Ray Kurtzweil, who was a student of mine, makes a... I could never find out his optical character reader worked.

NORBERG: Is that because he intended to manufacture it?

MINSKY: Yes, and he was protecting it and copyrighting it and keeping it a secret. It was really annoying. I mentioned the optical score reading. This is a kid named Alan Ruddenberg(?), he's in exactly the opposite situation, because MIT has just established a kind of proprietary rights agreement that students have to sign. And Ruddenberg wants to give the software away free. And they don't want him to.

NORBERG: That's interesting. That is a shift, certainly.

MINSKY: He's really furious. They're withholding his RA money until he signs the agreement.

NORBERG: Well, now getting back to this reviewing process of being on the phone continuously...

MINSKY: And e-mail. We were the first ones because of Larry Roberts. We simply use e-mail for ordinary social...

NORBERG: What I was trying to say, and trying to see if you would agree with me on this analysis, is that it seems to me that the present AI community, however we define what we mean here by community, the people who claim to belong to the artificial intelligence research community, seem to me to have a different attitude toward publication, for example, and how you evaluate publication, where many things appear in proceedings and preprints rather than in published articles and in refereed journals. This seems to arise from the kind of interaction you're talking about. That it's more important to interact with your colleagues in the same field than it is to worry about whether there is this trail behind you that you can point to as your productivity. Would you agree with that?

MINSKY: Well, there's both, because there's also the usual thing of people publishing incremental versions of their thesis...

TAPE 2/SIDE 1

MINSKY: I used to go out to Rand very often. I can't remember which years that was, but it was the years Newell and Simon were there and I think that's from the middle '50s; I'm not quite sure when it started.

NORBERG: What did you observe when you went out there?

MINSKY: Oh, I just hung around with my old graduate friends from Princeton, Lloyd Shapley and the game theorists were there. We used to do different little bits of mathematics and Newell and Simon were there and Cliff Shaw. So we would talk about problem-solving algorithms, but that's just a flashback.

NORBERG: But that's a fairly significant flashback, because in effect that is the AI community at the time. There were the people who were in the Cambridge area and then there are the people who were at Rand/Carnegie-Mellon.

MINSKY: That's right and I sort of.... Right. In fact, Rand was rich and all the other people used to come there. The guards knew us by face. We could just walk in and we could go out on the beach and talk about AI.

NORBERG: After the Mansfield Amendment or...

MINSKY: Or whatever it was that...

NORBERG: I'm going to accept your hypothesis about that, I'm still not convinced that that's correct.

MINSKY: I'm not either, but that's what we were told by the...

NORBERG: There is certainly an increase in bureaucratization within the Defense Department in those years; whether we want to ascribe it to the Mansfield Amendment or some other phenomenon, I think is irrelevant. But, that having happened, what happened to the activities here at MIT that you were involved in?

MINSKY: It didn't change an awful lot except that Seymour and I gave up the directorship and I think Pat [Winston] continued the lab in very much the same spirit of trying to get good, original people. But in fact he got very little research done. He didn't go back to research until about '78 or '79.

NORBERG: How did the lab grow in all those years? The number of people you've described, or at least by implication described, that were around in say 1960, seems to me to be about a dozen.

MINSKY: At most.

NORBERG: How did the size change over the next twenty years?

MINSKY: I think... I don't know exactly, I could look in the progress reports, but I think when Project MAC started then there was a sort of steady growth and by '73 there were probably 50 people, 40 or 50, and continued to grow. So now there's probably about 120, and it's at least two and a half floors. But also there's a discontinuous point where LCS absorbed first the MAC lab group and Jerry Sussman. The boundary between the AI Lab and LCS is indistinguishable now, and the whole thing is very big.

NORBERG: When Project MAC started in '63, did you people sever your association with RLE and the Computation Center?

MINSKY: Yes. Pretty much.

NORBERG: Because there are memos that continued to come out over the years after 1963 that are still ascribed to RLE and the Computation Center. It did fall off by '65.

MINSKY: I'm trying to remember. Tech Square must have started... The MAC summer thing... There was a big planning thing the first year and I think that was in Tech Square. That means that building is almost 20 years old.

NORBERG: Twenty or 30? 30 years old.

MINSKY: I think it was just finished when we moved in. So the connection with the RLE and the Computation Center just tapered off as people became more and more resident there. I think that may be more an architectural phenomenon than a deliberate management thing.

NORBERG: And the AI activity is still going on in Tech Square as I remember, on the 8th and 9th floors.

MINSKY: Yes. And 7th I'm happy to report. (laugh) Papert started to develop LOGO around the end of the '60s, '68-'69. So, that started another...

NORBERG: When you stopped being associated with directing the AI Laboratory, how did that change your activities?

MINSKY: I spent a lot of time building hardware for a while, but generally it didn't change it much because there hadn't been any bureaucracy before that and directing was a very small load of occasional decisions, mostly because I had Russell Nofsiger (?) there as General Manager. I hired him as an engineer and he started taking over things, pretty much managed the whole bunch of offices. They decided who would be in what rooms unless there was a complaint and then I would try to negotiate it. I'll bet I didn't manage more than four hours a week.

NORBERG: I would like to shift to a slightly different area here and that is the influence of AI on itself, first of all as a growing enterprise, intellectually; secondly its affect on computer science as the overarching discipline, and its affect on computing. I'd like to ask that in the following way. If you were me trying to do a study of AI that had an objective to demonstrate its development and its influence, I can't spend time on all of the sub-areas of AI, that just doesn't seem to be possible, there's too much going on.

MINSKY: So you have to do it by anecdotes or something.

NORBERG: Yes. What areas would you think have the most promise for demonstrating development and influence?

MINSKY: And again there are sort of things that you can still identify as AI and things that just happened to start developing in AI. A good example of the latter is modern complexity theory. When you take a course in computer science now, or if you go to graduate school, there are sort of four areas, and one of them is called algorithm theory. Algorithm theory is actually not how the algorithms work so much, but how you judge how efficient they are. I was very interested in the question of why does it take longer to compute some things than others. I got MIT to hire Albert Meyer who was very good at this. He sort of headed that section. So, in a very indirect way, it was the AI people who were... You might think it would be the business people who would say "What's the fastest possible way to sort." But as far as I can tell, the important thing in business is how can I get a sorter that's slightly faster than the competitors'. I don't want to discover one that's very fast, or else we'll all lose money. (Laugh) It's sort of funny. But that's what we're going through. So that's an important direction in theoretical computer science, that came from other places too. Shmuel Winograd at IBM was interested in that. They developed a pretty good [unintelligible]. What's another example....? That's a kind of weak example of a by-product. What's the question again?

NORBERG: The question is what sub-areas of AI would be the best ones to focus on to show development on the one hand and influence on the other.

MINSKY: Maybe robotics. As far as I can remember, modern robotics came almost entirely out of AI groups, but maybe that's a wrong judgement. I know people in Pennsylvania, at Ohio State or somewhere were trying to make a walking machine and it crept along. Nothing much came of it, but the idea of a robot with a vision system so it could see something and grab it... I don't think that happened anywhere else. We started doing it. SRI started doing it at Stanford, then somewhat later Carnegie-Mellon. I think in the next century that's going to be a very big thing. Right now it's impressive, but in fact, it hasn't dominated in the field. In industry you never lost a part once you've got it. It's always clamped or tied to a belt. So vision systems are used more for inspection than for assemblage. But generally the incentive for making more versatile robots has always come from AI. If you go to the robot store, the only thing you can buy are these things.

NORBERG: These things meaning hands? Clampers.

MINSKY: But if you look in the office next to mine you see Ken Stein (?) has a pretty good robot. If you go to Hitachi or somewhere you'll find that they've got some, too, for playing pianos and stuff, but I'd say most of the advanced robotics...

NORBERG: Can robotics be broken down itself into, say, vision and motion?

MINSKY: Yes. And the motion has been drifting out of AI and is being studied by people in other parts of the community, control theory people.

NORBERG: So is it just the vision that would be in AI now?

MINSKY: The vision is almost entirely in AI, and the planning. So between the vision and the motion where the machine must decide what should I pick up next, that's all in the AI places. Of course, it's not very important in factories where things are usually designed to be assembled in a certain way.

NORBERG: Okay, so that's two examples. Can we...

MINSKY: Speech recognition is a good example. That was done partly at Bell Labs for 40 or 50 years and never got anywhere, and a little bit of it was done in AI labs, mostly the Raj Reddy stuff at CMU. But Reddy got started because he was a student of McCarthy and he was interested in speech, and I persuaded him to see if he could do it without using Spectra (?), that was a very important _____________. He did a lot novel things that turned out... They ended up using Spectra, anyway...

NORBERG: Is that a relatively late development in comparison to robotics? Say in the '70s?

MINSKY: Speech? As far as I can remember Reddy's was the first one. Then there was a five-year ARPA program to push a speech program, which didn't pay off much. It _________ a couple of generations of speech machines. Yes, I'd say that's a little pocket and someone decided to put a lot of money into it, and it might have been a little too early. Speech is coming along now, but it seems that a lot of it is due to computers being a hundred times faster and algorithms only a little bit better. That's debateable. For example, OCR - Optical Character Recognition, a lot of AI people worked on that and a lot of inventors worked on it in different places. For example, Pat Winston's doctoral thesis was on it. Warren Titleman's thesis. He was one of the people from BBN. And now you can buy OCR machines that are actually pretty good. I tried one yesterday, put a full page in there and it only made a couple of errors; it was a magazine article with several pictures. It even got the captions to the pictures and put them in a separate file, it didn't confuse them with the text.

NORBERG: How about natural language processing?

MINSKY: That's been... Well, that certainly... This is parochial. I would certainly say that the AI people have done more on that than the linguists. Linguists tend to study syntax and avoid the question of meaning. But in general... So, the commercial natural language systems were mostly versions from people trained in AI labs, and they're not very good yet. But within a restricted domain they're pretty good. My feeling is that another 20 years is necessary because you can't deal with language unless you have the conceptual representation of what the words mean, and that's a big job for one project to try to do anything like that.

NORBERG: What's the basic problem there, can you tell me that? We have meaning.

MINSKY: Well, that's the point. To get the meaning, you have to solve a lot of the old problems of philosophy. For example, Douglas Lennart (?) in his psych project in Austin is trying to build a big database that has conceptual representations of all the sorts of things every child knows. It's a terrible job. [unintelligible]

NORBERG: Is it an accumulative task or does, every once in a while, the researcher come up against some new conceptual understanding that means you go back and you have to redo the whole thing again?

MINSKY: Very good. Excellent. It was supposed to be an accumulative task and they had to revise it twice so far and now they've got quite a few people entering knowledge and I feel that they made a big mistake at a certain point and they're going to have to do it over again. But Lennart (?) thinks it's good enough, so that when you have to do it over again it will know enough to help. I hope he's right.

NORBERG: (Laugh) I hope he's right, too.

MINSKY: See it knows a lot about things and classes of things. It knows that dogs are kinds of mammals and it also knows that animals have to eat and stuff like that. What it doesn't know is in the functional domain. It doesn't know that chairs are for people to sit in. And it's very important, I claim, that you know that, because to understand why a chair has legs, in fact why isn't a box a chair? The reason is you can't get...

NORBERG: You can't get your feet underneath it.

MINSKY: You can't stand... you can barely tolerate the couch. The chair is better and you can't tell Psych that yet, because it doesn't understand reason. Either the structure of the person... A really very close match. The real purpose of the chair is the seat floating in the air at the right height and the legs are not a proper part of the chair, they're just a heuristic for keeping the seat there, and he doesn't have a good way to express that. In recent years I've come to feel that this structure/function thing permeates everything. That there's no point in a thing unless it's for some purpose and the Psych thing wasn't made... So, I think the purpose of a thing, how this structure is related to someone else's intention, has to be almost a primitive thing there. Now he's got to jam these things in on top of all the other stuff. But we'll see. He says his logician knows how to do it. (Laugh) His logician is a young Indian mathematician named Guha, who is now a graduate student at Stanford. The MCC had to make a satellite office for Psych in Palo Alto.

NORBERG: I would have probably raised the ante a bit and besides vision and speech as two good examples of AI development, I guess I would have chosen something like problem-solving and knowledge based-systems as the third area.

MINSKY: Oh! The general development of expert systems is certainly the most productive of all. The problem solving has not developed into an industry exactly because I think it lacks this common-sense knowledge base and the whole world has to wait for things like Psych to work so that the machine just knows something about what you want done and what it's for. But the expert system... that's a very interesting thing because expert systems, you know in the commercial sense, these are these rule-based knowledge-base systems. The knowledge is put in the form of procedural rules which say if you want this then do that. They're not quite like GPS because they don't use differences. But what's important about that is that with one of these expert system shells, a person can write a problem solver by hand and you might be able to do something in a few days that would take a regular programming house a year to do. So, what it's done is it's wiping out conventional programming. Instead of writing this big, tricky program that has procedures that call other procedures in symbols you can't understand, you have this set of rules which say what it should do and then it just reads the rules. So it's sort of a computer; it's like you didn't need a compiler. I think what this is doing is, very soon this will be more than half of all software, so it's permeating everything. In fact, the ironic thing is that in the early '80s, basically Feigenbaum was the one who declared that rule-based systems were good enough for practical use. _____ members were telling him you're out of your head. They are just going to be always breaking; they're not rugged enough. But he was sort of right, and they actually started some companies because Feigenbaum invented the idea also, I think, that there was a special profession called knowledge engineers who would interview experts and formulate their decision rules and use things so they got investors to buy these companies. And I'm sort of amused to see that the reputation on Wall Street is rather bad. They say AI companies are losers; they cite symbolics. Of course symbolics didn't even do that kind of software. But although the AI business you can invest in didn't do well, if you look at a company like Arthur Andersen which half of its income comes from programming in this way and every big company has... Basically the only thing Feigenbaum might have been wrong about is that the knowledge engineers are not such a highly developed skill and someone else can learn that in a fairly short time. So you're right. The expert systems probably have the largest impact on the computer industry, because I really think they're going to eat up most of the programming jobs. It's just so much easier than writing a program. It's not very good if you want to write an actual low-level program that converts a certain file into something that could go out on a certain printer. You could write that as a rule based system, but it would be too slow.

NORBERG: Are there other areas in AI that I should consider? You have vision, speech, robotics generally, problem-solving, expert systems generally.

MINSKY: Expert systems is the biggest one. Well, there are these little things... Like we're trying to make some things that compose music. It's on knowledge-based ideas and things like that that might or might not work.

NORBERG: And they're all relatively new, aren't they?

MINSKY: Yes. And that sounds sort of frivolous except that there's something very peculiar about the entertainment industry, namely, it's the largest of all industries actually. So, who knows what would happen.

NORBERG: I notice you didn't mention graphics. Would you consider that AI development?

MINSKY: I'd say a lot of it is and a lot isn't.

NORBERG: It seems to me that you did some work, didn't you, that would now be considered graphics in trying to develop boundary understanding?

MINSKY: Well, the vision was a little bit related to the graphics. I did some of the first entertainment graphic things at BBN. They've been lost.

NORBERG: What were they like?

MINSKY: They made beautiful chaotic patterns of different sorts, making things out of curves instead of lines. Because all the early computer graphics was pretty much straight lines. But that's interesting. I think graphics is this case where the AI people did a fair amount of it. Gosper discovered these fractal distributions that are called Mandelbrodt distributions. Everybody in the AI community used _________ and things like that. At MIT the best graphics came from civil engineering. A guy named Charley Miller got interactive graphics going.

NORBERG: That wouldn't be considered AI, I assume.

MINSKY: I wouldn't consider it that. I would say that they didn't do that much direction.

NORBERG: Okay, so graphical techniques would not necessarily be included here.

MINSKY: Yes. You could ask... Ivan Sutherland's Sketchpad is the best graphic design program ever used, but it was never replicated. It's a miracle that's still being rediscovered. Have you talked to Alan Kay?

NORBERG: Just the meeting we had with him a week ago; I had a private chat with him on the side.

MINSKY: We have this permanent state of wonder that there'll be another, commercial version of Sketchpad. And some company up in Burlington made one recently but they ran out of budget so it's not available.

NORBERG: Why do you think it's not available?

MINSKY: I don't understand...

NORBERG: If I understand Sketchpad correctly, a number of the things that are included in the program, such as the window technique that Sutherland included in order to pull out sections of the bridge design that he was working on. If you think of the programs that now use that as a regular feature, then why wouldn't you say that Sketchpad had a substantial enough influence on development?

MINSKY: I think it probably did. People saw it, and they invented their own. The critical thing in Sketchpad was that it had internal constraints; so you could say this angle is twice as big as that one; as far as I know, no drawing program has that yet. It's quite easy to implement. It has this problem: if you have a set of constraints, then you have to have a program that solves them simultaneously and that could be hard. In fact, the way Sutherland did it was to take each of the errors and square them so it would be positive, and do a relaxation and that leads to [unintelligible]. It's hard to make one... I draw my lecture slides on MacDraw or something like that, and I'm always furious because I really want this thing to be twice that size at all times and I want to be able to move this... It's very funny that nobody's done that.

NORBERG: I wonder why.

MINSKY: Beats me. Actually what happens... No. I don't know why.

TAPE 2/SIDE 2

MINSKY: So you could say it inspired a lot of graphics.

NORBERG: Well, it's certainly seen as the starting point in almost anybody's account of what happened in computer graphics.

MINSKY: I'd say it's interesting that people haven't seen that the wordprocessor... For example, the first spell program, some kid at Stanford wrote the spelling program and he attached it to all of our editors and the whole environment of wordprocessing comes from AI. It's a big industry that's never called AI.

NORBERG: I'd like to go back to something you said earlier before getting on the second half of that question. There are two things that you said earlier, one of them will be the second half. Going back to the statement you made about AI being out there on the frontier continuously...

MINSKY: Well, I meant the place called an AI lab collecting people who...

NORBERG: Yes. Indeed. I accept that correction.

MINSKY: I don't say AIs are the only ones; they are the only ones I can think of. But I know that when I came to MIT, Gerry Wiesner in RLE had decided he would have McCullough. McCullough wasn't a professor, he didn't want to be; he had been a professor at Illinois. And he kept getting people. He had a guy named Manuel Serillo who was a Mexican conceptual artist. Serillo was making different theories of perception, and no one understood him very well. But Gerry got him; this might be a new deal. So when he was running RLE, he would collect people that he thought might be important for the future. It may not be inherent in artificial... I think it is inherent in the nature of artificial intelligence. Sure, you're trying to make machines that can think so you need ideas about thinking. But I hired composers who seemed to have... In Project MAC I hired three different composers at different times in the hope that they would, I guess they would seem more articulate than other musicians. They didn't pan out.

NORBERG: Was that to do computer music or to do things generally?

MINSKY: Yes. I was hoping they could do computer music and bring some new viewpoint. What happened was usually very ironic. Alan Fort was a composer, who's now head of the music department at Yale, and he wrote a book on SNOBOL programming. He got so interested in computers that he didn't, in fact, contribute anything to us, but that's the way it goes.

NORBERG: Sure. I want to go back to something you said before though when you talked about electronic mail.

MINSKY: In fact, I was furious because he won a prize for the best textbook on something that year. (Laugh) And I hoped he would write some new theory of musical expression.

NORBERG: I want to go back to the comment you made about E-Mail, when you said that in the reviewing process E-Mail became very significant as a replacement for, I presume, a replacement for the telephone and regular mail, postal service. How did it affect the interaction among people? And when did that affect begin?

MINSKY: Good question. And when was it realized? Because, I remember when people were talking about building the network. I think the idea in most people's minds was that if there was a program somewhere else, we could use it. See, the network was going to be a computational network. It is a bit, but mostly we used the network to get software from somewhere else, if you want, as a mail system, not as a computer. But, I would say, as soon as the thing started, we started sending mail to one another, and the glorious thing was the ARPANET because when Danny Bobrow went to Xerox... I don't remember whether the mail was working yet, when did Taylor go to Xerox?

NORBERG: About '71.

MINSKY: So the network was in place by then?

NORBERG: Yes, in '69.

MINSKY: So when I had a question, I'd E-Mail to Bobrow and just ask him and it was as though he'd never left. It's still that way over at electronics...

NORBERG: But how does that differ from the telephone? I mean, could you ask the same question over the telephone or is it more complicated?

MINSKY: Yes, but they wouldn't be there. The difference is that E-Mail is always answered. Usually within a day or so. But with a telephone they're not there and they say they'll call you back and they don't and it's very different. Also, it's more concise. You write the things and say, no I meant this. It's very nice. Then there's a sort of illegal aspect which is that it's used for all sorts of social things - the bootleg lists of the good Chinese Restaurants. One day I just noticed that I had gotten messages from all three children; Henry was teaching a course in LISP in Japan, Margaret was somewhere else, and Judy was somewhere else. It was this little family affair we had on the E-Mail and keep in touch that way. I probably am better connected. Until the last few years. When I started writing *Society of Mind*, I sort of stayed home and just worked, which was good, I didn't pay any attention to the world. But I bet that up to that point I was in closer touch with my students than a typical professor. No matter where they went they... You know, they wouldn't go somewhere that wasn't on the net.

NORBERG: Were there negative effects besides overloading the system with all sorts of social messages? Were there negative effects to this? Can you get too many messages that are really a problem to you?

MINSKY: Well, yes, but you make a filter. So some people complain. I don't make a filter usually. But, if you type control H, it gives you this list of the messages and then you can type D and it deletes one and it goes space and it skips. Or else you can tell it to put it in a box. I get a lot of messages about space affairs and I don't normally read them when they come, I just put them in a box and then if there's too many I erase the whole batch. (Laugh) It does take a few minutes.

NORBERG: I'm going to keep that in mind. Okay, my last question is back to this question of the influence of the developments in AI. You've mentioned several times now about the people down the hall who are in conventional areas of computer science. A number of them have never come around. They just don't believe that AI is a really significant part of the field, I guess.

MINSKY: Or they don't like something you can't prove.

NORBERG: Right. But it seems to me that there must be a number of examples that you could cite that are important contributions to their activities, now, non-AI activities, as a result of the work in AI. What are a few examples of that sort of contribution? Time-sharing is one you mentioned before.

MINSKY: Yes, time-sharing, and wordprocessing is probably the most important of all. I'm sure word-processing would have happened once time-sharing happened and certainly it would have happened once personal computers happened. So there's some inventions that are inevitable. Arcade games: a big prototype was the Space Wars that we had and, in fact, at some point we had to ban it, because it was so popular. The ban on Space Wars must have been about 1965 and it wasn't until 1975 or '80 that personal computers were cheap enough that you could have it again. So there's an invention that would have been made in 1975, if Steve Russell made it in 1960 that's just a sort of historical anomaly. So, you could say that the arcade games, computer games, came out of AI and the adventure games came out of the network for some reason, once people had time-sharing. But that's a kind of indirect thing.

Very direct tools are things like Mathematica which people buy for their computers, and applied mathematicians have these nice tools.

NORBERG: Are there no specific tools that were developed that people began to use then?

MINSKY: Well, MACSYMA and Mathematica are certainly the most clear cut ones because they actually came out of theses.

NORBERG: How about affects on programming generally?

MINSKY: Well, the real-time editors came out of AI. That is, in the old days you'd have a program and you'd run it and you'd stop the computer, load the program, and turn it on. I think we had the first one of those where you could stop the program or leave it running and make a window that would show the code being executed and edited right there. That led to the LISP machines, which were very popular until the Suns came along and the Sun machine would do the same thing; they were more mass produced. So the technology of how people make programs, certainly came from AI.

NORBERG: Yes. I guess that's the sort of conclusion I was reaching for: that there is an effect on the way programming is done as a result of AI development.

MINSKY: Yes. In fact, C has just become the most popular programming language in the world - in UNIX - and UNIX has not got this yet. We consider UNIX to be a worldwide disaster setting things back a few years.

NORBERG: You said that the other day and I wasn't quite sure what you meant.

MINSKY: It's very primitive. It puts you back in the early 1960s. Everyone tells me that in four or five years though they'll get back to the Project MAC way of doing things, it's just... The designer of UNIX didn't understand the importance of having a nice transparent program.

NORBERG: Is that possibly related to the problems that Bell had in continuing to be part of the MULTICS project?

MINSKY: That's a funny question.

NORBERG: Yes. It just occurred to me. It wasn't part of this design today.

MINSKY: Well there is this little fragment of history which is that John Pierce didn't like AI in any disguise. But he's been gone a long time and Arnold Penzias is the director now, and he likes the idea of AI. I don't think he's gotten around to following up on it, but I'm not sure.

NORBERG: Anything in data structures? Do these people develop databases differently now because of AI?

MINSKY: Well, almost everybody uses frames for representing ordinary knowledge, when they do ordinary knowledge. That's a popular thing. Some of the expert systems are critical. But for big commercial databases it doesn't pay to use anything fancy because you [unintelligible]. I think that the right person to ask about that would be someone who knows, like... Is Randy Davis on your list anywhere?

NORBERG: No.

MINSKY: He might be a good person to talk to because he came from AI and he's sort of the world's leading consultant on commercial expert systems.

NORBERG: Is he here in Cambridge?

MINSKY: Yes. In fact, the EE department wouldn't give him tenure, but the Sloan School did. Pretty scandalous. (Laugh)

NORBERG: I guess!

MINSKY: I wanted him to stay. In fact he's still in the AI lab as co-director.

NORBERG: Who else besides Davis would be interesting to talk to?

MINSKY: Well, Lennart would be pretty marvelous if you get a chance. I'm sure he has an interesting, different perspective. Roger Shank is another, a very good one.

NORBERG: You made a comment to Bill Aspray when he talked to you informally before about this project over a year ago now, in which you said that there were difficulties in getting funding out of place like the National Science Foundation in contrast to DARPA and IPTO.

MINSKY: Well, of course, maybe my point of view is so unique it can't be generalized, because we've just been through this thing where Wiesner gave us some money and then ARPA gave us lots of money for ten years. I did apply to NIH for some grant on sensory perception and they said, "No, we're not interested in that now." And I applied to NASA for something on telepresence (?) and they said "Well, no, we're interested in man in space." Every time I wrote a proposal to someone else, I never actually went all the way to getting a full-blown proposal. I said, "Boy it must be hard to do science these days, I'm so lucky." (Laugh) But I think it's not just me. I think if you talk to all the other people who were doing AI in the sixties with ARPA support they'll say the same thing.

NORBERG: If ARPA support is so readily available, then why would one bother to go for the other money anyway?

MINSKY: Well, one wouldn't. But you can see that it would be a very hard thing to do and, in fact, the AI lab got cutbacks from ARPA in the last three or four years and I don't think they've been able to replace it.

NORBERG: Last three or four years meaning in the late '80s?

MINSKY: Yes.

NORBERG: Do you remember the attempt to cut it back in the middle seventies?

MINSKY: Well, I think that must have been about when Pat took my job.

NORBERG: I see. Because Heilmeier apparently did think that AI was not something they should spend so much money on.

MINSKY: Well, that's when I quit, basically. It was him. But the guy before him was also hostile (?) to it and I couldn't understand this and then he went to TI and became the ...

NORBERG: ... AI guru there. That is quite remarkable.

MINSKY: But I never had any idea whether he was doing this on his own initiative or he was trying to rebalance ARPA because of some other cuts.

NORBERG: Were there some major meetings in which the AI people participated to convince him otherwise that you went to?

MINSKY: I never went to one. I think there were some minor meetings. I think he came to MIT once and I saw him in Pat's office.

NORBERG: Did you ever attend the PI meetings that DARPA ran?

MINSKY: A couple of them. They were great. I didn't go to them after I stopped being director, but I should have. I know some of my students went to them, even though they weren't PIs, and they said that was the best scientific meeting they'd ever been at. I heard that remark a lot of times.

NORBERG: Do you remember when that was?

MINSKY: No. Well, it must have been the... When did they start?

NORBERG: Well, let's see... The first ones that I remember are '65. But they also separated PIs and graduate students as well and held two different meetings.

MINSKY: Maybe I've got it muddled.

NORBERG: And the graduate students went there at a time when the PIs were told not to come so the graduate students could interact without the inhibition of having senior people around.

MINSKY: What a great idea.

NORBERG: That was Taylor's idea, apparently. And it worked and then it stopped and no one could resurrect it again. Now the PI meetings seem to be three, four, six hundred people at a time.

MINSKY: No kidding? Oh, they're not restricted to AI?

NORBERG: No. Well, there may be some sub-meetings with AI people, but the PI meetings, generally, have these large crowds of contractors, and it's very hard to do anything, I should think.

MINSKY: Then there was a period when they had these directed programs like autonomous vehicles. The speech one seemed like an all right idea. I think I was still directing then and I decided we wouldn't be a part of that.

NORBERG: Do you think the AI field is now too big?

MINSKY: Sure. In general, but the field of making these common sense databases is too small. I think they have to synchronize the various areas that are too small.

NORBERG: Would you like to separate them out from AI and call what's now become AI something else? Is that now part of conventional computer science, I guess is what I'm asking you.

MINSKY: You mean a lot of AI work?

NORBERG: Yes.

MINSKY: Yes, I guess so. I think expert systems activity is not returning very much scientific knowledge. [unintelligible] But you know, you've got this problem when you decide to separate something which is... and you can't decide... You don't want to get rid of the money cow, because you can hide basic research in large projects. But if they're out in the open, then they are vulnerable. And that, again, was a strange feature of that golden age of the ten years after Project MAC where the whole thing was full of basic research. I'll bet if we hadn't done those things like networks, wordprocessing, and time-sharing and so forth... It turned out that basic research actually produced a lot of maybe more useful stuff per person than applied research, but it's very hard to judge that now.

NORBERG: Well, if we look at the contributions of DARPA in terms of the things that they funded, it would seem to me that we would probably argue that most of the things that people cite as being the significant contributions tend to be out of basic research or at least a technique to stimulate basic research.

MINSKY: Yes. Well, what else did IPTO fund? I know ARPA is a billion dollar enterprise, so most of it I never heard of.

NORBERG: Well, if you think about the autonomous land vehicle which you mentioned in the strategic computing initiative. They label it as applied AI, so let's talk about it in the same way, I guess. Is that basic research? Trying to take an AI technique and make it do something for this machine - it's like an expert system.

MINSKY: Well, I think we'd have to read between the lines. It would be nice to get a pile of final reports from the different autonomous vehicle things and then you could see what fraction of it was just attaching things that people could do, and what fraction was a new vision system, new pattern system. I don't have that information, but I bet that at least half is genuinely basic, significantly basic research.

NORBERG: I guess I would say 30% rather than half. I think it's considerably less than half.

MINSKY: I haven't seen the figures. But even 30% is positive. If you took the speech thing, you'd probably have almost all basic research, but it was probably before its time because it was funded too enthusiastically and there weren't enough good ideas then. But if the whole thing cost 20 million, maybe the stuff that Reddy's group did was worth 20 million.

NORBERG: How would you characterize the contributions at DARPA? We're talking about just computer science.

MINSKY: Well, I told you in the elevator or whenever it was, my impression was that it was probably equal to all of NSF, because of very uniformed... It's just that all the good people I know are in the DARPA circuit. Every now and then I meet somebody who isn't in our network and I'm sort of surprised - I didn't know you could do good work out there. (laugh)

[INTERRUPTION]

MINSKY: Another million dollars we could order somebody to get all those reports
and add up how many.

NORBERG: I want to ask you one off-the-wall question here as a way of bringing this to, what I hope is an adequate close. That is, you've been listening to me now for almost two hours ask various questions, some of which are coherent, some of which are not, in terms of the overall questioning process. Are there surprises that you didn't hear? That is, is there something you didn't hear that you would consider to be a surprise in my questioning you.

MINSKY: Some question you should have asked?

NORBERG: Right.

MINSKY: I can't think of one.

[INTERRUPTION]

MINSKY: I can't think of any. Well, the general question, is there something they could've done better? I don't know what it is.

NORBERG: No, I guess we can't do that. See one of the things that I'm trying to keep away from here is to talk about success. It seems to me that's not what an historical analysis of their activity should focus on.

MINSKY: Oh no, I meant is there some way of funding that they should've done better.

NORBERG: I certainly understood your comment, but I think what most of the people who have listened to the DARPA crowd in Washington want to say is that somehow that program was unique. What I'm trying to find out is what about the program would contribute to uniqueness. Seems to me uniqueness should be a conclusion rather than an hypothesis to be demonstrated.

MINSKY: Well, I don't know if it's unique, but I'd attribute the great success and influence to the initial heads of that office, namely Lick and was it Roberts or Sutherland?

NORBERG: Sutherland next, then Taylor, then Roberts.

MINSKY: Taylor was before Roberts?

NORBERG: Yes. Lick brought him in.

MINSKY: But anyway there was something miraculous about that particular string of people. I didn't know about Taylor much, but I know the other three were sort of in this position of molding the thing, and really had a lot of good taste about what were the important problems to look at and that's very rare.

NORBERG: How would they know the right problems to look at, do you think?

MINSKY: I guess we'd all work together. (laugh) My problems are the ones that I think are the right problems. (laugh)

NORBERG: All right. So, that means that convergence between the research community and the people in those offices handing out the money is an important element.

MINSKY: Right. And the easiest way is if one of the people will actually volunteer to go there and help. And for some reason they got that many of them and you usually can't get any.

NORBERG: Up through Kahn that certainly worked.

MINSKY: That's right. Now if Kahn hadn't -- I don't know how Kahn did so well. But he did and he hadn't been in it very much before then.

NORBERG: Well, he was another BBN person who went down there.

MINSKY: Okay. So getting a great director of an office was... And then I that Marvin Denicoff's personal influence also sort of roaming around the projects and sometimes deciding what was going well. I don't know what he did about it, but...

NORBERG: That would be from the ONR side.

MINSKY: That's right. That was a piece of luck, having a contract administrator actually know about the personal lives of these people in a quiet, sympathetic way. I have a feeling he must have told some of them they were wasting their time. Who knows?

NORBERG: Did you see a difference between the programs of places like NSF and NIH, in computing now, and a place like DARPA/IPTO?

MINSKY: Yes. On my few visits to NSF I would meet people who were trying to do a good job, but they didn't know what was going on. It was a kind of review process in which you either have to have a peer review process or a benevolent dictatorship.

END OF INTERVIEW